arXiv:1511.00725v2 [cs.LG] 25 Jan 2016

Galaxy-X: A Novel Approach for Multi-class Classification in an Open Universe[☆]

Wajdi Dhifli, Abdoulaye Baniré Diallo[1]

Department of Computer Science, University of Quebec At Montreal, PO box 8888, Downtown station, Montreal (Quebec) Canada, H3C 3P8

Abstract

Classification is a fundamental task in machine learning and artificial intelligence. Existing classification methods are designed to classify unknown instances within a set of previously known classes that are seen in training. Such classification takes the form of prediction within a closed-set. However, a more realistic scenario that fits the ground truth of real world applications is to consider the possibility of encountering instances that do not belong to any of the classes that are seen in training, *i.e.*, an open-set classification. In such situation, existing closed-set classification methods will assign a training label to these instances resulting in a misclassification. In this paper, we introduce *Galaxy-X*, a novel multi-class classification method for open-set problem. For each class of the training set, Galaxy-X creates a minimum bounding hyper-sphere that encompasses the distribution of the class by enclosing all of its instances. In such manner, our method is able to distinguish instances resembling previously seen classes from those that are of unseen classes. To adequately evaluate open-set classification, we introduce a novel evaluation procedure. Experimental results on benchmark datasets as well as on synthetic datasets show the efficiency of our approach in classifying novel instances from known as well as unknown classes.

Keywords: Multi-class classification, Open-set classification, Galaxy-X

1. Introduction

Classification is a central task in machine learning and many other related fields including data mining, pattern recognition and artificial intelligence. It has an extremely large number of domains of application ranging from aid decision systems in finance and marketing, to bioinformatics and computer vision [1][2]. The most conventional classification scenario is to train a classifier on

[☆]Implementation and experimental datasets are available on https://sites.google.com/site/wajdidhifli/softwares/galaxy-x.

Email address: dhifli.wajdi@courrier.uqam.ca, diallo.abdoulaye@uqam.ca (Wajdi Dhifli, Abdoulaye Baniré Diallo)

a set of instances of known classes, *i.e.* the training set, then to predict the class label of unknown instances within the same set of already seen classes [3][4][5][6]. Such classification takes the form of prediction within a *closed-set* of labels. However, in many real world applications, due to the growth of data collection, the training data could only represent a partial view of the domain and thus may not contain training exemplars for all possible classes. In such scenario, the classifier may be confronted, during prediction, to observations that do not belong to any of the classes seen in training. This makes the classification become within an *open-set* of labels where the presence of observations from unseen classes is possible. In *open-set classification*, traditional closed-set classifiers will fail in the prediction for observations of unseen labels.

In applications where the user is interested in the identification of few classes from a large classification universe, the most conventional way is to fuse the set of uninteresting classes into one single negative set which usually makes the dataset becomes highly unbalanced. In this case, the classifier becomes overwhelmed by negative observations which hinder the discrimination of positive classes. Some attempts have emerged trying to remedy such situation, mainly based on the sampling of a subset of representatives from the negatives [7]. However, it is very difficult and somehow unfair to reduce all the negatives into a small summary that may not be sufficient to represent the whole set. A more appropriate transformation of such problem is an open-set classification where only positive classes are modeled in training and any observation that remarkably deviates from the distribution of known classes is rejected.

Several domains of application of open-set classification exist. In the following, we present bioinformatics and computer vision examples. In bioinformatics [2], the advances in sequencing technology have made the acquisition of genomic sequences fast and easy. While a virologist analyses the genomic sequence of a virus, he always keeps the possibility that the latter can be an unknown virus. A closed-set classification does not help in such situation since the new virus will be assigned a previously known type (label) preventing a discovery. Another example of domains of application is computer vision [8]. For instance, in face identification, the system is only interested in recognizing a number of faces within an infinite set of possibilities, *i.e.* an open universe of classes. In similar applications, the classifier should be able to create a decision boundary that envelops the class instances and resembles its distribution such that whatever lies outside of the class boundary is rejected.

In this paper, we introduce *Galaxy-X*, an open-set multi-class classification approach. For each class, Galaxy-X creates a minimum bounding hyper-sphere that encloses all of its instances. In such manner, it is able to distinguish between novel instances that fit the distribution of known classes from those that diverge from it. Galaxy-X introduces a softening parameter for the adjustment of the minimum bounding hyper-spheres to add more generalization or specialization to the classification models. To properly evaluate open-set classification, we also propose a novel evaluation technique, namely *Leave-P-Class-Out-Cross-Validation*. Experimental evaluation on simulated as well as benchmark datasets shows the efficiency of Galaxy-X in open-set multi-class classification.

2. Related Work

Very few works have addressed open-set classification in the literature. Scheirer et al. presented a formalization of open-set classification and showed its importance in real world applications [8]. They discussed the bias related to the evaluation of learning approaches on popular datasets. They showed how recognition accuracies are inflated in closed-set scenarios, leading to an over-estimated confidence in the evaluated approaches [8, 9]. In binary classification, SVM defines an hyper-plane that best separates between two classes. Scheirer et al. proposed an SVM based open-set multi-class classifier termed one-vs.-set SVM [8], which defines an additional hyper-plane for each class such that the latter becomes delimited by two hyper-planes in feature space. A testing instance is then classified as of one training class or as of an unknown class, depending on its projection in feature space. Although this strategy delimits each training class from two sides, the class "acceptance" space is left unlimited within the region between the hyper-planes and no additional separator is provided to prevent misclassifying unknown instances that lie within the same hyper-planes bound but far away from the training class distribution in feature space.

Semi-supervised classification [10] have addressed open-set classification to some extent, where part of the dataset is unlabeled and the goal is to label as much as possible of the unlabeled data then use them in training to enhance the performance. Unlabeled instances are labeled based on their distance from the distribution of labeled data where far instances can be rejected by all classes. Although this somehow resembles open-set classification, in this context the acceptance/rejection is performed in training and the goal is to minimize the loss on the training set to optimize the classification performance. However, in open-set classification, the acceptance/rejection is performed in prediction on testing data to classify them as of one training-class or as of an unknown one.

Another important learning approach for open-set problems is one-class classification. The most known technique is one-class SVM [5] where the classifier is trained only on a single positive class and the aim is to define a contour that encloses it from the rest of the classification universe. Any instance that lies outside of the defined class boundary is considered as negative. One-class classification is mainly used in outlier and novelty detection. It is limited to single class classification and cannot be directly used in multi-class classification.

One-vs.-one and one-vs.-rest [6] are popular techniques for multi-class classification. One-vs.-one constructs a model for each pair of classes and test examples are evaluated against all the models. A voting scheme is applied and the predicted label is the one with the highest number of votes. One-vs.-rest creates a single classifier per class, with the examples of that class as positives and all the other examples as negatives. In prediction, all classifiers are applied on the test example and the predicted label is the one with the highest confidence score. It is possible to use one-vs.-rest for open-set classification by iteratively considering each class as the positive training set, and all the remaining (known) classes as the rest of the classification universe. However, in open-set classification, the rest of the classification universe is (at least theoretically) unlimited

and thus the classifier will suffer a negative set bias.

Based on [11] and [12], it is possible to build a simple open-set multi-class classifier using a combination of a one-class classifier and a multi-class classifier. In the first step, all training classes are fused into a single large super-class and the one-class classifier is trained on the entire super-class. In this setting, the one-class classifier testing instances that do not fit the distribution of all known training classes are directly rejected and labeled as unknown. In the second step, the multi-class classification is trained on the original training classes and used to classify instances that were not rejected by the one-class classifier.

3. Galaxy-X

3.1. Preliminaries and Problem Definition

Let $\mathcal{D}$ be a training set of n instances and $\mathcal{L}$ be the set of possible labels in $\mathcal{D}$, $\mathcal{D} = \{(x_1, l_1), ..., (x_n, l_n)\}$ where $l_i \in \mathcal{L}$ and x_i is defined by a vector in d-dimensional space, $\forall i \in [1, n]$. In open-set classification, the classifier should be able to assign to a test instance x a label l_x that is known $l_x \in \mathcal{L}$ or that is *unknown*, *i.e.*, $l_x \in \mathcal{L} \cup \{"unknown"\}$. In this setting, it is necessary to define a boundary envelop for each class in order to make it distinguishable from other unknown possibilities. The definition of such boundary is hard and delicate as the delimited-class-space should reflect the class distribution by enclosing as much as possible of its instances while keeping outside as much as possible the rest of instances. Indeed, this can be seen as an optimization problem of the classification error that considers a trade-off between generalization and specialization. As a possible solution, we define the minimum bounding hyper-sphere $\mathcal{M}$ as the smallest hyper-sphere that circumscribes all instances of a considered class. For a class $\mathcal{D}_l \subseteq \mathcal{D}$ of label $l \in \mathcal{L}$, the hyper-sphere $\mathcal{M}_l$ represents the class model that resembles the distribution of $\mathcal{D}_l$ instances. Each class model $\mathcal{M}_l$ ($\forall l \in \mathcal{L}$) is defined as:

$$\mathcal{M}_l = (c_l, r_l), \forall l \in \mathcal{L}, r_l > 0 \tag{1}$$

where c_l is the center of $\mathcal{M}_l$ hyper-sphere (the class mean $\overline{x}$):

$$c_l = \overline{x}, \forall x_i \in \mathcal{D}_l, \forall l \in \mathcal{L} \tag{2}$$

and r_l is the radius of $\mathcal{M}_l$ hyper-sphere, *i.e.*, the distance between c_l and the most divergent instance from $\mathcal{D}_l$ (the class variance):

$$r_l = \max\left(distance(x_i, c_l)\right), \forall x_i \in \mathcal{D}_l, \forall l \in \mathcal{L} \tag{3}$$

where *distance* is a function returning the distance between c_l and x_i with respect to a distance measure. In a multi-class classification scenario, the resulting representation space is similar to a galaxy of classes in an open universe of possibilities.

3.2. The Training Process

Algorithm 1 describes the training phase in Galaxy-X. Given a training set $\mathcal{D}$ and a training label set $\mathcal{L}$ over $\mathcal{D}$, we create a model $\mathcal{M}_l$ for each class $l \in \mathcal{L}$ that is composed of the class minimum bounding hyper-sphere center c_l and radius r_l as defined in Equations 1, 2 and 3.

Algorithm 1: Galaxy-X: The training process

Data: $\mathcal{D}$: training set, $\mathcal{L}$: training labels
Result: $\mathcal{M}$: set of class models

```
1 begin
2     M ← ∅
3     foreach (l ∈ L) do
4         c_l ← Centeroid (D_l)
5         r_l ← Boundary (D_l)
6         M_l ← (c_l, r_l)
7         M ← M ∪ M_l
```

3.3. Acceptance of Instances

In open-set classification, the classifier should be able to discriminate between instances of the different known classes and to reject those of unknown classes. Therefore, we define a score of acceptance of an instance by a class depending on its position from the class boundary.

Definition 1. *(Acceptance Score) The acceptance score for an instance x by a class of label $l \in \mathcal{L}$, is defined as follows:*

$$AcceptanceScore(x, l) = 1 - \frac{distance(x, c_l)}{r_l} \tag{4}$$

where distance is an appropriate distance measure, c_l is the center of the class of label l, and r_l is its radius.

The acceptance score is defined in $]-\infty, 1]$ ($i.e. AcceptanceScore \in \mathbb{R}_{\leq 1}$) and it allows to decide whether an instance is accepted or rejected by a class. The score is interpreted as follows:

- $AcceptanceScore \in [0, 1]$: the query instance x is accepted by the class l:
  - $AcceptanceScore \in]0, 1[$: x is inside the minimum bounding hyper-sphere of l,
  - $AcceptanceScore = 1$: x is in the class center, *i.e.*, $x = c_l$,
  - $AcceptanceScore = 0$: x is on the class boundary, *i.e.*, $distance(x, c_l) = r_l$.

- *AcceptanceScore* < 0: x is out of the class boundary (rejected). The lower is the *AcceptanceScore*, the farther is x from the class distribution.

Galaxy-X tries to minimize the classification error (Err) that can be formulated as:

$$Err = \sum_{\forall l \in \mathcal{L}} \sum_{\forall x_i \in \mathcal{D}} F(x, l) \tag{5}$$

where $F \in [0, 1]$ is a binary function that is defined as follows:

$$F(x,l) = \begin{cases} 1, & \text{if } \mathit{AcceptanceScore}(x,l) \in [0,1] \text{ and } x \text{ belongs to } l, \text{ or} \\ & \text{if } \mathit{AcceptanceScore}(x,l) < 0 \text{ and } x \text{ does not belongs to } l \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

3.4. The Classification Process

3.4.1. Filtering Prediction Candidate labels

Based on the acceptance score, it is possible, for a given query instance x, to filter the subset of candidate labels $\mathcal{L}_x$, $\mathcal{L}_x \subseteq \mathcal{L}$. The latter is the subset of remaining possible candidates, such that if $\mathcal{L}_x \neq \emptyset$, then the predict label l_x is an element of it, $l_x \in \mathcal{L}_x$. The general algorithm of filtering of the candidate class labels is described in Algorithm 2. It starts with an empty set of candidate labels. Given the set of training class models, it tests whether the query instance x is accepted or rejected by each training class according to Definition 1. Indeed, it rejects all class labels where the query instance do not fit the class distribution, *i.e.* when x lies outside of the class boundary. Only the subset of labels of accepted classes is retained as the possible candidate labels for prediction.

Algorithm 2: Galaxy-X: The label filtering process

```
Data: M: set of class models, x: test instance
Result: L_x: retained candidate labels
1 begin
2     L_x ← ∅
3     foreach (class model M_l ∈ M) do
4         if AcceptanceScore (x, l) ≥ 0 then
5             L_x ← L_x ∪ l
```

3.4.2. Handling Class Overlapping

It is possible to obtain a non intersecting set of minimum bounding hyper-spheres in the case where training classes are perfectly separable. In such case, if a query instance x is circumscribed by an hyper-sphere then x takes the latter's class label otherwise x is considered as of an unknown class. However, in real world cases the hyper-spheres may overlap mainly in the presence of high inter-class similarity. In fact, the overlapping space between classes resembles a local

closed-set classification within an open-set classification context. In this case, a local closed-set classifier is trained only on the overlapping classes then used only for classifying query instances that lie within the overlapping space, *i.e.*, instances that are accepted by multiple classes in Algorithm 2, $|\mathcal{L}_x| > 1$.

3.4.3. The Classification Process

Algorithm 3 describes the classification process of Galaxy-X. After training, the first step in prediction is the filtering of candidate labels according to Algorithm 2. If the retained set of candidate labels is an empty set $\mathcal{F}_x^L = \emptyset$, then the query instance x do not fit any of the training class distributions. Since it is an open-set classification, the predicted label l_x is set to *"Unknown"*. If $|\mathcal{F}_x^L| = 1$, x is only accepted by one training class. In this case, the predicted label is that single filtered possibility $l_x \leftarrow \mathcal{F}_x^L$. In the case where $|\mathcal{F}_x^L| > 1$, x shares a degree of similarity with more than one class and lies in the overlapping area between the hyper-spheres of the retained class labels. As this situation presents a conventional closed-set classification, a closed-set classifier $\mathcal{E}$ is locally trained only on the retained classes of $\mathcal{F}_x^L$, then $\mathcal{E}$ is used to predict the class label l_x of x such that $l_x \leftarrow \mathcal{E}(x)$ and $l_x \in \mathcal{F}_x^L$.

Algorithm 3: Galaxy-X: The classification process

Data: $\mathcal{M}$: set of class models, $\mathcal{E}$: local closed-set classifier, x: test instance
Result: l_x: predicted label for x

```
1  begin
2      F_x^L ← FilterLabels(M, x)
3      if F_x^L = ∅ then
4          l_x ← "Unknown class"
5      else
6          if |F_x^L| = 1 then
7              l_x ← F_x^L
8          else
9              Train(E, F_x^L)
10             l_x ← E(x)
```

3.5. Softening Class Boundaries

In order to add flexibility to the models, we introduce a softening parameter $\delta \in \mathbb{R}$ that allows to perform a distortion of the class boundary. Indeed, it allows to add more generalization or specialization to the classification models as a trade-off between sensitivity (recall) and specificity. Figure 1 shows respectively examples of positive and negative softening of a class boundary. In Figure 1(a) a positive softening extends the radius of the minimum bounding hyper-sphere allowing to add more generalization to the class model. Extending the

class boundary may help detecting test instances that are from the same class but slightly deviate from the training instances. In contrast, in Figure 1(b) a negative softening is performed shrinking the radius of the hyper-sphere adding more specialization to the class model. Shrinking the class boundary may help rejecting instances that do not belong to the class but are within the class hyper-sphere near to the class boundary. In addition, it can be used to alleviate or remove overlapping between classes. If performed, the softening value has to be carefully chosen as an over-generalization engenders many false positives. Whereas an over-specialization looses true positives and the model will only fit a small portion of the class instances.

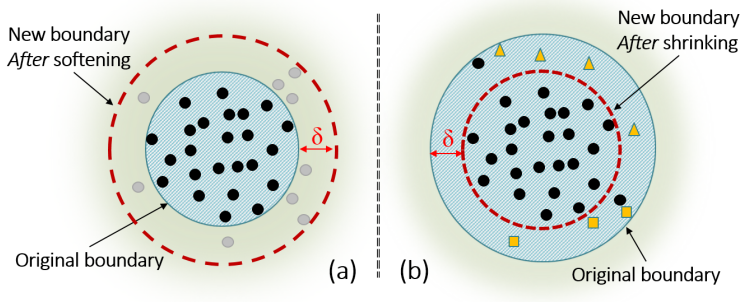

Figure 1: Example of positive (a) and negative (b) softening of the class boundary. Black dots represent training instances of the class. Gray dots are test instances from the class but slightly differ from its distribution. Squares and triangles are negative samples and δ is the softening magnitude. Softening the boundary can help detecting true positives (the gray dots) or rejecting false positives (the squares and triangles).

Definition 2. *(Soft Acceptance Score) The softening can be introduced in the acceptance score. We define the soft acceptance score as follows:*

$$SoftAcceptanceScore(x, l, \delta) = 1 - \frac{distance(x, c_l)}{r_l + \delta} \tag{7}$$

where δ is the softening parameter, distance is an appropriate distance measure, c_l is the center of the class of label l, and r_l is its radius.

Similarly to acceptance score, soft acceptance score is defined in $]-\infty, 1]$ and is interpreted in the same way. It is worth noting that softening can also be introduced in training (instead of the soft acceptance score) in the definition of class boundaries such that line 5 in Algorithm 1 becomes $r_l \leftarrow$ `Boundary` (D_l) $+ \delta$. According to Equations 5 and 7, the optimal δ value, denoted δ^*, should be the one that minimizes the classification error as follows:

$$\delta^* = \underset{\delta}{\operatorname{argmin}}\, Err = \underset{\delta}{\operatorname{argmin}} \sum_{\forall l \in \mathcal{L}} \sum_{\forall x_i \in \mathcal{D}} F(x, l, \delta) \tag{8}$$

where F is defined similarly to Equation 6 but based on the $SoftAcceptanceScore$.

Lemma 1. *Given a classification scenario $S_{\mathcal{D}}$, a classification performance evaluation technique Perf, and a closed-set classifier $\mathcal{X}$:*

$$\forall S_{\mathcal{D}}, Perf(Galaxy\text{-}\mathcal{X}, S_{\mathcal{D}}) \geq Perf(\mathcal{X}, S_{\mathcal{D}})$$

Proof. In the worst case, the optimal softening value will be very high until training models completely overlap resembling a closed-set classification. In this case, evaluation instances will be classified using the local closed-set classifier. Consequently, $Perf(\text{Galaxy-}\mathcal{X}, S_{\mathcal{D}}) = Perf(\mathcal{X}, S_{\mathcal{D}})$. □

4. Experimental Evaluation

Evaluating open-set multi-class learning methods requires defining proper measures and protocols.

4.1. How Open is an Open-set Classification?

We propose *Openness* as a measure to quantify the openness of a classification scenario ($S_{\mathcal{D}}$).

Definition 3. *(Openness) It measures the ratio of labels that are unseen in training but encountered in prediction from all labels of the dataset $\mathcal{D}$. Openness is defined as follows:*

$$openness(S_{\mathcal{D}}) = \frac{|UnseenLabels|}{|\mathcal{L}|} \tag{9}$$

Openness is defined in $\mathbb{R}^+$. An *openness* value of 0 means that it is a closed-set classification scenario, otherwise it is an open-set classification. Theoretically, the value of *openness* can be even $+\infty$ meaning an infinite set of possibilities. However, in practical cases, the number of test labels can usually be delimited. In our experiments, *openness* $\in [0, 1[$, since the open-set classification will be simulated from a benchmark dataset were all possible labels are known, *i.e.*, $|\mathcal{L}| = |TrainingLabels| + |UnseenLabels|$. An *openness* of 1 means that $|TrainingLabels| = 0$ corresponding to a clustering context which is out of the scope of this work.

4.2. Evaluation Technique

Careful experimental procedures need to be designed to appropriately evaluate multi-class open-set classification. Conventional evaluation techniques including hold out, cross validation, random sampling, and their variants, are not suitable for open-set classification. They were originally designed for closed-set classification and hence they do not present sufficient restrictions on labels to simulate an open-set classification evaluation. We propose *Leave-P-Class-Out-CrossValidation* a novel evaluation technique for open-set classification. It allows to simulate an open-set classification that better resembles real-world applications where we do not have knowledge of all classes in training. The general procedure of *Leave-P-Class-Out-CrossValidation* is described in Algorithm

4. First, all possible combinations $\mathcal{C}$ of $\mathcal{P}$ labels from $\mathcal{L}$ are computed. In each iteration, one possible combination *comb* is randomly chosen from $\mathcal{C}$, without replacement. All instances of a label l_{comb}, $\forall l_{comb} \in comb$, are temporarily discarded from the dataset $\mathcal{D}$ to be directly added to the test set. These instances are referred to as the *Leave-out-instances*. All labels in *comb* are unseen in training but encountered in testing which simulates open-set classification. An $\mathcal{N}$-fold-cross-validation is performed on the remaining instances, $\mathcal{D} \backslash$ *Leave-out-instances*, where in each cross validation *Leave-out-instances* are directly added to the test set. The evaluation is repeated until a maximum number of iterations α is reached or no more combination is possible.

Algorithm 4: Leave-P-Class-Out-CrossValidation

Data: $\mathcal{D}$: classification dataset, $\mathcal{L}$: the set of labels of $\mathcal{D}$, α: maximum number of iterations, $\mathcal{P}$: number of labels to leave out in each iteration, $\mathcal{N}$: number of cross validation folds, $\mathcal{E}$: the open-set classifier

Result: *Scores*: Average classification scores

```
1  begin
2      C ← All possible combinations of P labels form L
3      while (α > 0) and (C ≠ ∅) do
4          Randomly chose a combination comb from C
5          Leave-out-instances← all instances of D_{l_comb} | ∀l_comb ∈ comb,
           comb ⊆ L
6          foreach TrainingSet, TestSet ∈
           N-CrossValidation(D\Leave-out-instances) do
7              Train(E, TrainingSet)
8              TestSet ← TestSet ∪ Leave-out-instances
9              PredictedLabels ← Predict(E, TestSet)
10             Scores ← Scores ∪ Statistics(PredictedLabels)
11         C ← C\comb
12         α ← α − 1
```

4.3. Evaluation Measures

The natural way to evaluate classification is to use the accuracy measure which refers to the amount of correctly classified instances from the dataset. In multi-class classification the accuracy is averaged over all classes of the dataset. However, in open-set classification, the negative set can extremely outnumber the positive set which inflates the classification results causing an overestimation of the performance of the classifier. Moreover, the number of testing classes is (at least theoretically) undefined. F-measure (also so-called f-score), which is the harmonic mean of precision and recall, represents a good alternative

for open-set classification. Formally, it is defined as:

$$F\text{-}measure = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{10}$$

Where Precision is $\frac{TP}{TP+FN}$ and Recall is $\frac{TP}{TP+FP}$. We use the weighted version of f-measure as the evaluation metric for our experiments. F-measure is computed for each label, then results are averaged, weighted by the support of each label which makes it account for label imbalance.

4.4. Experimental Protocol and Settings

In order to guarantee an equal participation of all used attributes in the classification, a min-max normalization is applied on each attribute independently such that no attribute will dominate in the prediction ($x_{normalized} = \frac{x-min}{max-min}$, where x is an attribute value, min and max are the minimum and maximum values for the attribute vector). In each experiment, we use Galaxy-X to classify a considered dataset in a simulated open-set classification using the Leave-P-Class-Out-CrossValidation evaluation. The maximum number of iterations α is set to 10 and the number of cross validations in each iteration is 5. We evaluate the classification performance in terms of weighted f-measure using incremental values of *openness*. Classification results are compared with the gold standard multi-class classification strategy One-vs.-Rest [6] using a linear SVM as the baseline classifier (OvR-SVM), and with the open-set multi-class classifier One-vs.-Set SVM (OvS-SVM) proposed by Scheirer et al. [8]. OvS-SVM was used with the default parameters as requested by authors, where the generalization/specialization of the hyper-planes are performed automatically through an iterative greedy optimization of the classification risk. We also build a two-step open-set multi-class classifier, termed OCSVM+OvR-SVM, based on [11] and [12] as discussed in related work (Section2). In the first step of OCSVM+OvR-SVM, a one-class SVM (OCSVM) with an RBF kernel is trained on the entire training instances considered as a single super-class. For OCSVM, instances that strongly deviate from the super-class are labeled as "unknown". Otherwise, the instance is passed to the OvR-SVM for classification where OvR-SVM is trained on the original training classes using a linear SVM. For Galaxy-X, we use the same closed-set classifier as OvR-SVM, OvS-SVM, and OCSVM+OvR-SVM (*i.e.* SVM with a linear kernel). We show results of Galaxy-SVM using a fixed softening value δ=-0.3 (*i.e.* -30% expressed in terms of class radius). We also show results of H-Galaxy-SVM our approach (denoted shortly for Hyper Galaxy-SVM) using the optimal δ value for each openness where δ^* is obtained through a greedy search within a range of [-0.5, 0.5] with a step size of 0.1. The used distance measure for our approach is the euclidean distance. It is worth noting that Galaxy-X is not limited to SVM but other closed-set classification algorithms can be used as well. In contrast, OvS-SVM is limited to the SVM framework. Thus, we use SVM for Galaxy-X, OvR-SVM and OCSVM+OvR-SVM for consistency.

5. Results and Discussion

5.1. Evaluation on Classification of Handwriting Digits

We first evaluate our approach on a dataset of handwriting digits[1]. The dataset is composed of 1797 instances divided in 10 classes representing the Arabic digits. Each instance is an 8x8 image of a hand-written digit, and thus represented by a vector of 64 features of values between 0 and 16 respectively to the gray-scale color of the feature in the image. Figure 2 shows a visualization of the handwriting digits dataset. Figure 2a shows a sample from the dataset representing examples of each handwriting digit. As the dataset is multidimensional, we use manifold learning, a non-linear dimensionality reduction approach, to visualize the distribution of instances of the dataset. We use the t-distributed Stochastic Neighbor Embedding (t-SNE) [13] which converts similarities between data points to joint probabilities and tries to minimize the Kullback-Leibler divergence between the joint probabilities of the low-dimensional embedding and the high-dimensional data. Figure 2b shows a visualization of the handwriting digits dataset using t-SNE where each data point is colored according to ground truth class membership. We highlight the separability of the clusters distribution where each class can approximately be completely distinguished from the rest of the dataset.

(a) A selection from the handwriting digits[1] (b) T-SNE visualization

Figure 2: Visualization of the handwriting digits dataset. Colors in (b) are according to ground truth class membership.

Figure 3 shows f-measure results of Galaxy-SVM using different δ values in a simulated open-set classification of *openness*=0.5, meaning that only 5 classes are seen in training and all the 10 classes are encountered in prediction. The obtained results are compared with those of SVM (the used local closed-set classifier). The classification performance of SVM is very low as the classifier is at least incapable of correctly classifying the 5 classes that were unseen in training. SVM is originally designed for closed-set classification scenarios and thus it will assign one training label to all test instances of the 5 unknown classes resulting in a misclassification. Galaxy-SVM highly outperforms SVM

[1] http://scikit-learn.org/stable/auto_examples/datasets/plot_digits_last_image.html

in terms of f-measure in the best case. However, with higher values of softening, the performance of Galaxy-SVM leans toward that of SVM. This is due to the effect of over-generalization since the minimum bounding hyper-spheres become progressively larger with higher δ values until they completely overlap. In this setting, no rejection will be performed and only the local closed-set classifier of Galaxy-SVM will be used to classify all instances. With lower values of δ, the class boundaries become tighter adding more specialization to the class models. This allows Galaxy-SVM to better reject instances that do not resemble the overall distribution of the training classes. However, the value of δ should be carefully specified since an over-specialization leads to a high distortion of the models making them incapable of covering the variance of training classes. The value of δ^* is the one that guarantees the highest f-measure representing the best trade-off between generalization and specialization with respect to the classification scenario.

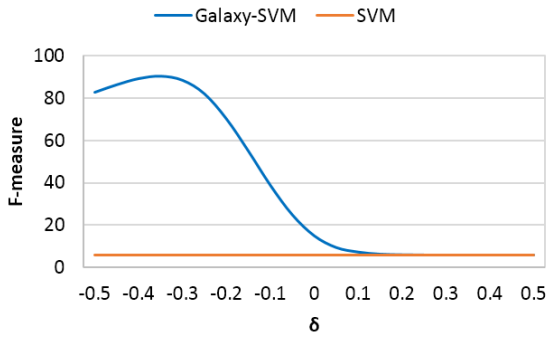

Figure 3: F-measure performance in open-set classification of the handwriting digits dataset for OvR-SVM and Galaxy-SVM with *openness* of 0.5 and using different δ (softening) values.

Figure 4a shows the classification performance in terms of f-measure for H-Galaxy-SVM (using the optimal δ value in each iteration), Galaxy-SVM (with a fixed δ = -0.3), OvS-SVM, OCSVM+OvR-SVM, and OvR-SVM using different *openness* values. The value of *openness* ranges from 0 to 0.8 corresponding to a number of held-out classes (P) from 0 to 8 that is used in the Leave-P-Class-Out-CrossValidation. An *openness* value of 0 corresponds to a closed-set classification meaning that all testing classes are seen in training. In this case, the classification performance of all classifiers are approximately the same since in the absence of rejected classes they all perform at least as good as the used closed-set classifier, *i.e.* SVM. *Openness* values from 0.1 to 0.8 correspond to open-set classification. The downward tendency of OvR-SVM is clear with higher *openness* values. In fact, the more open the classification is, the more false assignments OvR-SVM will generate. In contrast, all open-set classifiers maintained higher f-measure performance than OvR-SVM due to their ability to reject instances from unseen classes. Both H-Galaxy-SVM and Galaxy-SVM (δ=-0.3) outperformed all the other approaches in open-set classification sce-

narios. Using a fixed δ value of -0.3, Galaxy-SVM was able to give very close results to those of H-Galaxy-SVM meaning that $\delta^* \approx \delta$ in all cases. Overall, f-measure results of OCSVM+OvR-SVM were higher than those of OvS-SVM except for *openness*=0.2 where they gave similar results and for *openness*=0.1 where OvS-SVM outperformed OCSVM+OvR-SVM. Figure 4b shows rejection f-measure results on the held-out classes for each *openness* value. H-Galaxy-SVM and Galaxy-SVM scored better rejection f-measure than OvS-SVM and OCSVM+OvR-SVM in all cases providing the best trade-off between rejection-precision and rejection-recall. It is worth noting that although OCSVM+OvR-SVM and OvS-SVM uses the same closed-set classifier (*i.e.* SVM), in opposite to f-measure results for *openness*=0.1 and *openness*=0.2, OCSVM+OvR-SVM provided better rejection f-measure than OvS-SVM. A possible explanation for these results can be that in apposite to OCSVM+OvR-SVM, OvS-SVM performs an additional hyper-plane optimization for SVM. While OCSVM+OvR-SVM was more accurate in rejecting unknown instances than OvS-SVM, the latter provided a more accurate multi-class classification for the known classes in *openness*=0.1 and *openness*=0.2 settings.

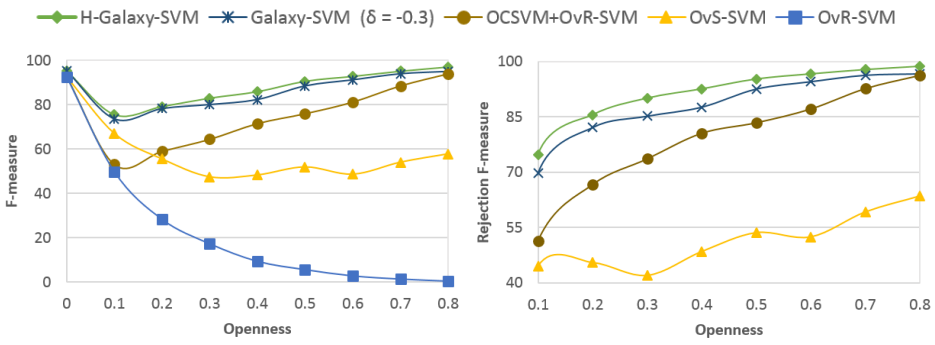

(a) (b)

Figure 4: F-measure (a) and rejection f-measure (b) results of H-Galaxy-SVM, Galaxy-SVM (δ=-0.3), OvS-SVM, OCSVM+OvR-SVM and OvR-SVM in open-set classification of the handwriting digits dataset with different *openness* values.

5.2. Evaluation on Face Recognition

We evaluate Galaxy-SVM on face recognition using the Olivetti faces dataset from AT&T Laboratories Cambridge[2]. This dataset consists on a set of 400 pictures, 10 pictures each of 40 individuals. The pictures were taken at different times, varying the lighting, facial expressions (open/closed eyes, smiling, *etc.*) and facial details (with/without glasses, *etc.*). Each picture is of size 64x64 resulting of a feature vector of 4096 values of gray levels. Figure 5a shows a

[2]http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html

sample from the dataset representing face picture examples for 8 individuals. The task is to identify the identity of the pictured individuals. Figure 5b shows a t-SNE visualization of the dataset, where each data point is colored according to ground truth class membership. In opposite to the handwriting digits dataset, we notice a high inter-class overlapping between the different clusters making it difficult to isolate each class separately. With so many classes, such inter-class overlapping, and only 10 examples per class, the classification of this dataset is very challenging. Transforming this dataset into an open-set classification makes the classification even more challenging.

(a) A selection from the Olivetti faces[2]

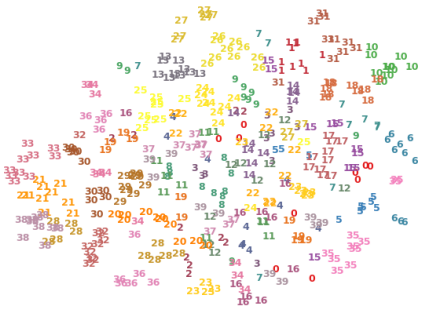

(b) T-SNE visualization

Figure 5: Visualization of the Olivetti faces dataset. Colors in (b) are according to ground truth class membership.

Figure 6 shows the classification results in f-measure using different values of δ at *openness*=0.5 meaning that 20 classes are seen in training and all the 40 classes are encountered in prediction. Again, SVM performance was very poor compared to that of Galaxy-SVM in the best case where our approach outperformed SVM by more than 80% order of magnitude. Indeed, even with no softening (δ=0), Galaxy-SVM was able to outperform SVM by approximately 60% order of magnitude.

Figure 7 shows the classification performance in terms of f-measure (Figure 7a) and rejection f-measure (Figure 7b) for H-Galaxy-SVM (using δ^* in each iteration), Galaxy-SVM (with a fixed δ=-0.3), OvS-SVM, OCSVM+OvR-SVM, and OvR-SVM using different *openness* values. The value of *openness* ranges from 0 to 0.8 corresponding to a number of held-out classes (P) from 0 to 32 with a step size of 4. As shown in the figure, Galaxy-SVM handles higher values of *openness* better than all the other approaches. Indeed, even at an extreme *openness* value of 0.8 corresponding to only 8 training classes and 40 testing classes comprising 32 classes that were unseen in training, Galaxy-SVM was able to classify known as well as unknown class instances with high f-measure of almost 95%. H-Galaxy-SVM outperformed all the other approaches in open-set classification cases. H-Galaxy-SVM and Galaxy-SVM (δ=-0.3) gave close results for open-set classification cases except for *openness*=0.1 where H-Galaxy-SVM performed better. This can be explained by the fact that in that case more general-

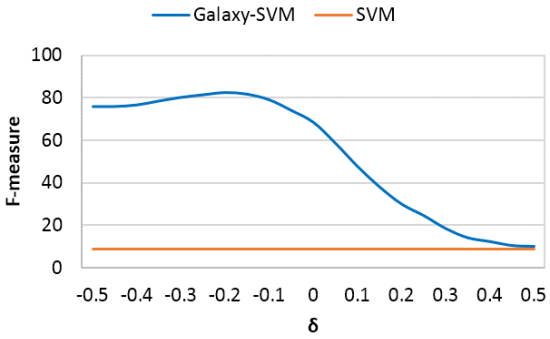

Figure 6: F-measure performance of Galaxy-SVM and OvR-SVM in open-set classification of the Olivetti faces dataset with *openness* of 0.5 and using different δ (softening) values.

ization was needed whereas Galaxy-SVM (δ=-0.3) performed a specialization of -0.3. This conclusion is supported by the f-measure result of the closed-set classifier OvR-SVM in that case, where it outperformed Galaxy-SVM (δ=-0.3) with no rejection at all. Even though H-Galaxy-SVM and OvS-SVM used the same closed-set classifier (SVM) and gave very similar results in terms of rejection f-measure, H-Galaxy-SVM outperformed OvS-SVM in terms of classification f-measure in all open-set classification cases. This is due to the difference between the class representation models used in each approach. Our approach encapsulates each class with a minimum bounding hyper-sphere that isolates it from the rest of the classification universe from all sides. However, OvS-SVM defines two hyper-planes for each class that delimit the latter from only two sides in feature space leaving the class "acceptance" space unlimited within the region between the hyper-planes as discussed in Section 2. The classification technique of our approach is more efficient in such classification scenarios with high inter-class overlapping. The classification performance of OCSVM+OvR-SVM compared to that of OvS-SVM is in contrast to that obtained with the handwriting digits dataset. Indeed OvS-SVM outperformed OCSVM+OvR-SVM in most open-set classification cases of the Olivetti faces dataset. This is due to the high inter-class overlapping that prevents the one class classifier OCSVM from efficiently isolating the training classes (when considered as one super-class) from the overlapping unknown classes. This is clearly illustrated in the rejection f-measure results in Figure 7b where OCSVM+OvR-SVM scored less than the other open-set classification methods.

5.3. Evaluation on Synthetic Datasets

We further evaluate the classification performance of H-Galaxy-SVM, Galaxy-SVM (δ=-0.3) on the classification of multiple synthetically generated datasets. We also compare the obtained results with those of OvS-SVM, OCSVM+OvR-SVM and OvR-SVM in open-set classification of the same datasets using different *openness* values ranging from 0 to 0.8. The datasets are composed of

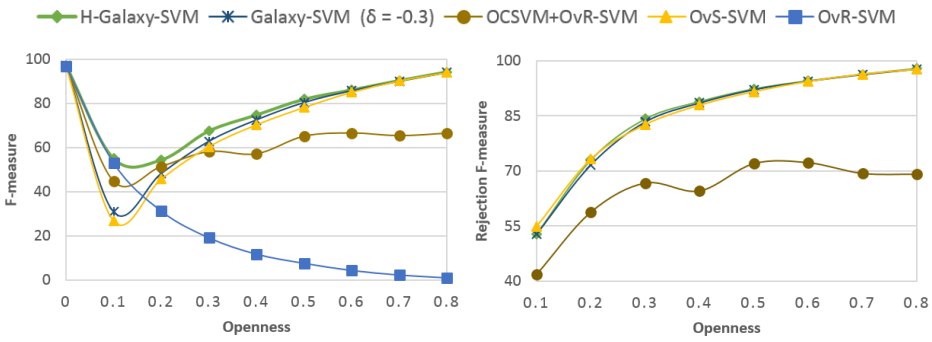

(a) (b)

Figure 7: F-measure (a) and rejection f-measure (b) results of H-Galaxy-SVM, Galaxy-SVM (δ=-0.3), OvS-SVM, OCSVM+OvR-SVM and OvR-SVM in open-set classification of the olivetti dataset with different *openness* values.

different numbers of classes ranging from 10 to 50 classes. For all the datasets, each class is Gaussian distributed and is composed of 500 instances such that in each cross validation evaluation 100 instances of each class will be used in testing and 400 instances in training. Each instance is composed of a two dimensional numerical vector and the standard deviation for each class distribution is 1. The number of repetition of the experiments is 10 at each setting. This means that for each dataset and at each *openness* value the total number of classification evaluations is 500 (number of iterations α=10 * 5 cross-validation * 10 simulations). In order to make the classification harder, we delimit the classification space for the synthetic datasets by introducing the following two constraints:

- for each dataset, the distance between the center (the mean) of each class and the center of any other class from the dataset should be no less a minimum distance of 2, and
- the distance between the center (the mean) of each class and the center of at least one other class from the dataset should not be greater than 5.

The Figures 8a, 8b, 8c, 8d and 8e show the f-measure results for all the classification methods using respectively 10, 20, 30, 40, and 50 classes at different values of *openness*. Overall, the five figures show similar tendency for all the classification approaches. Indeed, in accordance with previous results on real datasets, the best f-measure results are obtained using our approach (accordingly H-Galaxy-SVM and Galaxy-SVM (δ=-0.3)) while OvR-SVM gave the lowest results with a significant decrease in classification performance at higher values of *openness*. OvS-SVM and OCSVM+OvR-SVM showed better f-measure results than OvR-SVM in open-set classification cases, however, the classification performance of our approach was far better than both OvS-SVM

and OCSVM+OvR-SVM in all cases. To better show the effect of increasing the number of classes on the classification performance of all the classification methods, we show the obtained f-measure results at *openness*=0.5 for different number of classes ranging from 10 to 50 classes. The Figure 8f shows the obtained f-measure results. It is noticeable that overall, the classification performance of OvS-SVM, OCSVM+OvR-SVM and OvR-SVM are slightly decreasing with higher number of classes. In fact, the f-measure results for OvS-SVM, OCSVM+OvR-SVM and OvR-SVM are respectively 32.86%, 57.75%, and 6.16% for a setting of 5 training classes and 10 evaluation classes. For the same approaches, the results are respectively 23.94%, 24.56%, and 2.59%. In contrast, increasing the number of classes did not significantly impacted the classification performance of our approach. Indeed, the obtained f-measure results for H-Galaxy-SVM, Galaxy-SVM (δ=-0.3) are respectively 92% and 84.81% for the setting of 5 training classes and 10 evaluation classes, and 91.28% and 85.78% for the setting of 25 training classes and 50 evaluation classes.

6. Conclusion

In this paper, we addressed a fundamental problem in machine learning and artificial intelligence namely open-set classification where it is possible to encounter, in prediction, instances of classes that were unseen in training which better fits many real world applications. In open-set classification, it is necessary to define a decision boundary for each class that envelops its instances and resembles its distribution. The definition of such class boundary is difficult as it should offer the best trade-off between generalization to detect test instances that belong to a training class but slightly deviate from its distribution, and specialization to reject unknown instances that do not resemble any known class distribution. In many real world applications where the closed-world hypothesis does not hold, it is important to detect such unknown instances and raise the attention of experts to address them separately, preventing a misclassification. We introduced Galaxy-X, an approach for open-set multi-class classification. Galaxy-X encapsulates each class with a minimum bounding hyper-sphere that resembles the class distribution by enclosing all of its instances. In such manner, our method is able to distinguish instances resembling previously seen classes from those that are unknown. Galaxy-X presents high flexibility by a softening parameter that allows extending or shrinking class boundaries adding more generalization or specialization to the classification models. Experimental results on classification of handwriting digits and face recognition show the efficiency of Galaxy-X in open-set classification compared to gold standard approaches from the literature. An evaluation procedure was also introduced to adequately evaluate open-set classification.

An interesting future work is to propose a representation model for non spherical like shaped classes in order to avoid the risk of over-generalization in empty regions of the hyper-sphere. Furthermore, another direction is to develop an estimation method for a fast discovery of the optimal (or near optimal) δ, avoiding a greedy search across all possibilities. A preliminary idea is to

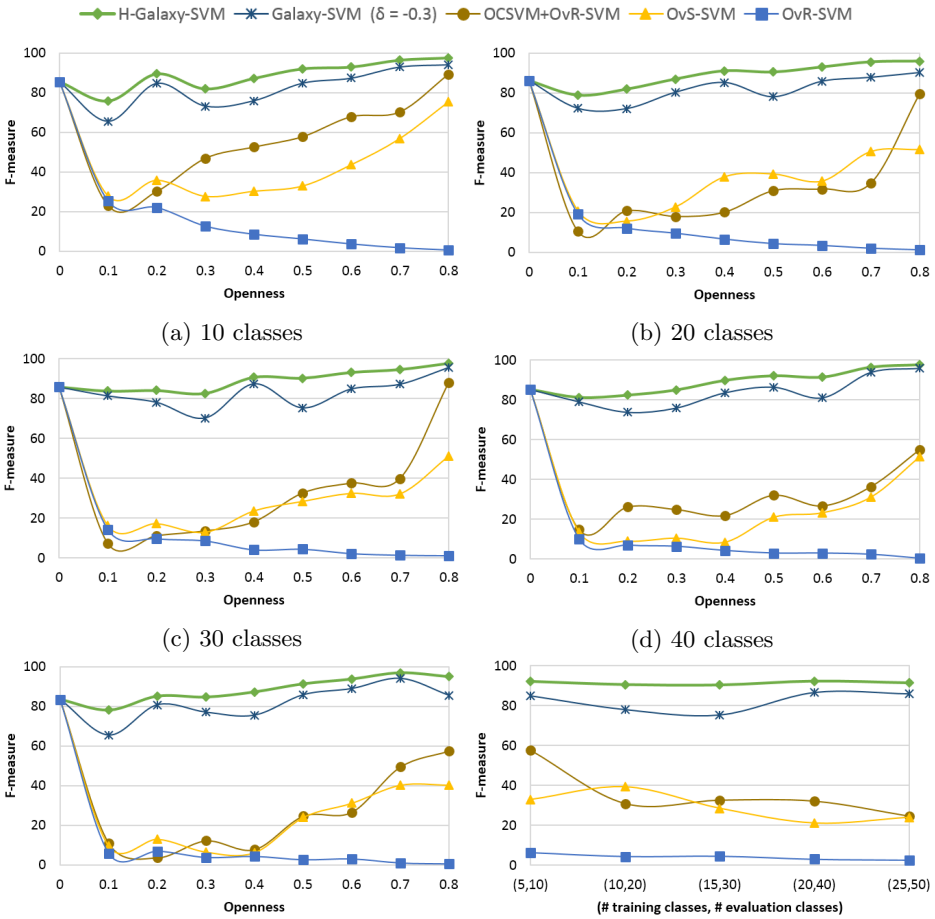

(a) 10 classes

(b) 20 classes

(c) 30 classes

(d) 40 classes

(e) 50 classes

(f) Different numbers of classes

Figure 8: F-measure results of our approach (H-Galaxy-SVM and Galaxy-SVM (δ=-0.3)) compared to OvS-SVM, OCSVM+OvR-SVM and OvR-SVM in open-set classification of multiple synthetically generated datasets. The classification results in 8a, 8b, 8c, 8d and 8e are depending on the number of classes and the *openness* value. Results in 8f are depending on the number of classes at *openness*=0.5.

examine the classification error on instances that are close (to some extent) to class boundaries with different δ values. The classification error will only be approximated on this small portion of instances allowing a quick estimation of optimal δ. The search ends when an optimal performance is reached with respect to an evaluation measure.

References

References

[1] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: 26th Annual Conference on Neural Information Processing Systems (NIPS'12), Curran Associates, Inc., 2012, pp. 1097–1105.

[2] M. W. Libbrecht, W. S. Noble, Machine learning applications in genetics and genomics, Nature Reviews Genetics 16 (2015) 321–332.

[3] J. R. Quinlan, Simplifying decision trees, International Journal of Man-Machine Studies 27 (3) (1987) 221–234.

[4] C. Cortes, V. Vapnik, Support-vector networks, Machine Learning 20 (3) (1995) 273–297.

[5] D. Tax, One-class classification: Concept learning in the absence of counter-examples, Ph.D. thesis, Technische Universiteit Delft (2001).

[6] A. Rocha, S. Goldenstein, Multiclass from binary: Expanding one-versus-all, one-versus-one and ecoc-based approaches, IEEE Transactions on Neural Networks and Learning Systems 25 (2) (2014) 289–302.

[7] B. Raskutti, A. Kowalczyk, Extreme re-balancing for svms: A case study, ACM SIGKDD Explorations Newsletter 6 (1) (2004) 60–69.

[8] W. Scheirer, A. Rocha, A. Sapkota, T. Boult, Toward open set recognition, IEEE Transactions on Pattern Analysis and Machine Intelligence 35 (7) (2013) 1757–1772.

[9] A. Torralba, A. A. Efros, Unbiased look at dataset bias, in: Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'11), IEEE Computer Society, 2011, pp. 1521–1528.

[10] X. Zhu, A. B. Goldberg, R. Brachman, T. Dietterich, Introduction to Semi-Supervised Learning, Morgan and Claypool Publishers, 2009.

[11] T. Landgrebe, P. Paclík, D. M. J. Tax, R. P. W. Duin, Optimising two-stage recognition systems, in: Proceedings of the 6th International Conference on Multiple Classifier Systems, MCS'05, Springer-Verlag, 2005, pp. 206–215.

[12] D. M. J. Tax, R. P. W. Duin, Growing a multi-class classifier with a reject option, Pattern Recognition Letters 29 (10) (2008) 1565–1570.

[13] L. Van Der Maaten, Accelerating t-sne using tree-based algorithms, Journal of Machine Learning Research 15 (1) (2014) 3221–3245.